THE ECONOMIC VALUE OF TEMPERATURE FORECASTS IN ELECTRICITY GENERATION

BY THOMAS J. TEISBERG, RODNEY F. WEIHER, AND ALIREZA KHOTANZAD

United States electricity generators save $166 million annually using 24-h temperature forecasts to improve the mix of generating units that are available to meet electricity demand.

Weather forecast services in the United States are mostly provided to users free of charge by the government. The economic justification for this approach is that weather forecasts belong to a class of goods that economists call "public goods." These are goods that would typically be undersupplied and underused in the absence of public provision.[1] However, there are many competing uses for public funds, and it is appropriate to ask whether value is being received for funds expended on weather forecasts. It is also appropriate to ask whether the incremental value from improved forecasting capabilities justifies the extra cost in providing these capabilities. The present study provides some of the important information needed to answer these kinds of questions.

This study estimates the cost savings (i.e., benefits or value) attributable to temperature forecasts used by the U.S. electricity generation industry in planning how to produce electricity up to 24 h ahead in time. The focus is on temperature because it is the key weather variable affecting the demand for electricity, particularly in regions of the country where there is heavy use of air conditioning.[2] Accurate temperature forecasts can improve the accuracy of electricity demand forecasts, and better demand forecasts can lower electricity production costs. This is true because electricity is typically produced by a variety of generating units with different lead times to be readied for service, costs of being readied for service, and production costs once in service. Having the best mix of generating units available at the right time saves money.

The paper begins with a discussion of relevant results from the earlier work by Hobbs et al. (1999),

[1] See Gunasekera (2004) for an elaboration of these ideas as they apply to weather forecasts.

[2] Other aspects of weather that may affect electricity demand, though to a much lesser extent, include clouds, wind, and humidity.

AFFILIATIONS: TEISBERG—Teisberg Associates, Charlottesville, Virginia; WEIHER—National Oceanic and Atmospheric Administration, Silver Spring, Maryland; KHOTANZAD—Pattern Recognition Technologies, Inc., Dallas, Texas
CORRESPONDING AUTHOR: Thomas J. Teisberg, 1475 Ingleside Drive, Charlottesville, VA 22901
E-mail: tjteisberg@compuserve.com
DOI:10.1175/BAMS-86-12-1765

In final form 22 July 2005

which estimates the savings in electricity generation costs from improvements in electricity demand forecasts. Next, our estimates of the relationship between the quality of temperature forecasts and the quality of electricity demand forecasts is presented for four different temperature forecasts at six sites around the United States. Using these results in conjunction with the Hobbs et al. (1999) work makes it possible to estimate the cost savings from different temperature forecasts at these sites. These site-specific cost savings are then extrapolated to estimate total benefits, and incremental benefits, for the United States as a whole. A concluding section presents some further discussion of our results.

COST SAVINGS FROM IMPROVED ELECTRICITY DEMAND (OR LOAD) FORECASTS.

The costs of generating electricity are lower with a good forecast of electricity demand, or electricity load, as it is sometimes called. Such cost savings can be understood as the avoided costs of errors that result from inaccurate forecasts. For example, it is often true that generating units with shorter production lead times have higher production costs. Thus, if a load forecast is too low, long lead-time units will not have been made ready, and the more expensive short lead-time units will be pressed into service to supply the unanticipated demand. Alternatively, if a load forecast is too high, long lead-time units may be readied for production unnecessarily, and the costs of doing so are incurred unnecessarily. These are only a couple of the possible scenarios that produce higher costs if load forecasts are inaccurate.

In the short run, higher costs would reduce the generating company's operating profits. In the longer run, however, higher production costs tend to produce higher electricity prices for consumers, due to competition, which is an increasingly important factor in the electricity industry. Regardless of who bears the burden of increased costs, they represent a net loss to the U.S. economy insofar as resources are expended unnecessarily to supply electricity demand.

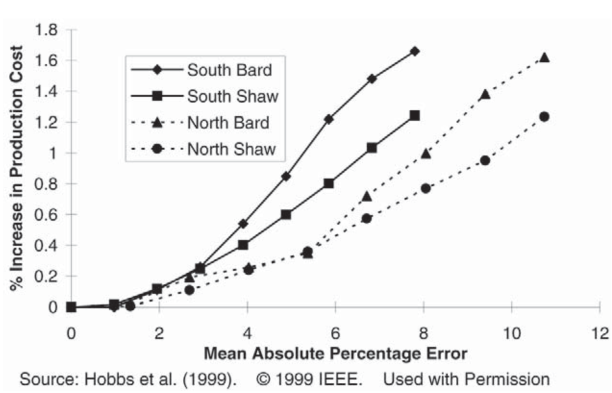

Fig. 1. Load estimation errors increase costs.

A study by Hobbs et al. (1999) estimates cost savings in electricity generation from better forecasts of electricity loads. In the Hobbs et al. study, actual loads and forecasted loads were obtained for two utility systems—one northeastern and one southern. These forecasted loads were produced by a model developed by the Electric Power Research Institute and used by a number of electricity generators (Khotanzad et al. 1995, 1998). Greater or lesser forecasting accuracy was simulated by scaling the actual forecast errors (up or down), thereby creating alternative forecasts with identical distributional characteristics, but differing accuracies.

These simulated load forecasts were then evaluated in the context of two alternative specifications of generating systems that might be employed to meet the electricity demand. These systems are referred to as Bard and Shaw, after the primary authors of the published papers from which the generating system specifications were obtained (Bard 1988; Shaw 1995, respectively). The result of the simulation is four relationships between the simulated forecast error and the economic cost associated with that error—one relationship for each load forecast (North and South) and one for each generating system (Bard and Shaw). These four relationships are shown in Fig. 1, which is reproduced from Hobbs et al. (1999) with permission.[3]

The cost increases due to load forecast inaccuracy range from about 0.35% to 0.85%, for example, when the mean absolute percentage error (MAPE) of the load forecast is 5%. At most values of MAPE, this cost increase is significantly higher for the southern load and for the Bard generating system. Of course, the cost of forecast inaccuracy for either region and generating system is lower when MAPE is lower (and it is zero if MAPE is zero, i.e., the forecast is perfect). Hobbs et al. (1999) note that for a system with a typical \$20 per megawatt hour (MWh)-averaged production cost, the annual losses from forecast inaccuracy would range from about \$600 to \$1,500 per megawatt (MW).[4]

[3] This is Hobbs et al.'s (1999) original Fig. 3, which has been very slightly modified for clarity of exposition.

[4] Hobbs et al.'s (1999) calculations here are percentage savings (0.35%–0.85%) times annual cost per megawatt. Annual cost per megawatt is the average cost per megawatt hour (\$20) times the annual hours of operation (8760).

TEMPERATURE FORECASTS AND DEMAND FORECASTS.

To use the results in the Hobbs et al. (1999) study to assign value to the temperature forecasts used in electric power generation, it is necessary to estimate the relationship between the temperature forecast accuracy and load forecast accuracy. This is done using a proprietary neural network–based load forecaster provided by Pattern Recognition Technologies, Inc. (Dallas, Texas). This load forecaster, like the model used in the Hobbs et al. study, is initially "trained" using data on actual loads and other variables that are predictive of loads. In our analysis, these other variables were the previous day's load, temperature, and temperature forecast for the current day. After training, the model is used to predict next day loads using the predictive variables as inputs.

Our analysis requires data on actual electric loads and corresponding location-specific temperature observations and forecasts for a number of locations in the United States. Such data are available, but only for a limited number of sites. From the available alternatives, six were chosen; these were located in Vermont, Ohio, Florida, Texas, southern California, and Washington. This set of sites provides reasonable geographic coverage of the continental United States and encompasses a wide range of differing climatic conditions. Because electric utility companies consider the electric load data to be proprietary, the specific power systems used and the raw data itself cannot be disclosed.

The temperature forecast data for the chosen sites were obtained from DTN Online (www.dtnonline.com is a Web-based provider of a wide variety of data, including weather forecast data). These data consist of temperature forecasts obtained from a National Weather Service (NWS) computer model known as the model output statistics (MOS) aviation guidance (MAV; see Dallavalle et al. 2004). MAV includes forecasts, at 3-h intervals, for temperature and other weather variables for specific geographic locations. These forecasts are provided both to NWS for use in constructing official forecasts and to private sector meteorologists.[5]

To investigate the implications of better-or-worse temperature forecast accuracy, three additional forecasts are constructed. Two of these, the persistence and perfect forecasts, are constructed to span the full spectrum of forecast quality. The persistence forecast uses today's actual temperature as a forecast for tomorrow's temperature; it is a naïve forecast that requires no skill. The perfect forecast uses tomorrow's actual temperature as today's forecast for tomorrow, and obviously represents perfect skill in forecasting.

The third forecast constructed is intended to represent the official NWS forecast. This forecast is the raw MAV, improved slightly to reflect the typical improvements made by human forecasters when they adjust the raw MAV to arrive at official NWS forecasts. To construct this forecast, recent forecaster improvements over MAV, in the next day's maximum temperature forecast, are averaged for the 7 A.M. and 7 P.M. (local time) forecasts of the preceding day. The average forecaster improvements are 7.0% (NWS Eastern Region), 10.5% (NWS Southern Region), and 11.4% (NWS Western Region). These improvements are applied, respectively, to MAVs for the Vermont and Ohio sites, the Florida and Texas sites, and the southern California and Washington sites, to construct the NWS forecasts.[6]

Using approximately 2 yr of data available for each of the six sites in the study, the load forecasting model was trained on the first year's data, and then run for the second year's data to obtain the results used in this study. Table 1 shows the load forecast error results for the six sites. For comparison, the load forecast errors reported by Hobbs et al. (1999) from their northeastern and southern sites are also reported.

There are a couple of things to note in Table 1. First, the percentage reduction in MAPE for the NWS forecast relative to the persistence forecast is very similar for the four nonwestern sites, ranging from 27% (Ohio) to 35% (Florida and Texas). However, error reduction is notably smaller for the western sites (2% for southern California and 13% for Washington). These observations suggest that 24-h

[5] Strictly speaking, the NWS does not call MAV computer model outputs a "forecast," but instead refers to them as "guidance." This is to make clear the distinction between these numbers and the official forecasts of the NWS. However, MAV outputs are NWS products that are made available for use by the public, including electric utilities that use them for forecasting future electric loads. Thus, this paper takes the liberty of loosely referring to the MAV outputs as forecasts, even though they are formally called guidance.

[6] Our data show considerable variation across local NWS offices in forecaster improvements over the MAV forecasts, and it is also possible that there is considerable variation over time for a given office. Because the objective here is to estimate the value of the temperature forecasts on a national basis, the regional-averaged forecaster improvement rates are used.

Table 1. MAPE for six sites and alternative weather forecasts.

	Persistence	MAV	NWS	Perfect
Vermont	2.94	2.04	2.01	1.78
Ohio	4.58	3.40	3.36	2.56
Florida	5.43	3.65	3.48	2.54
Texas	6.14	4.19	3.98	2.66
Southern California	3.13	3.31*	3.08	1.82
Washington	2.83	2.56	2.47	2.14
Hobbs et al.'s Northeast			5.4	
Hobbs et al.'s South			3.9	

*Here the MAV produces a larger MAPE than does the persistence forecast. But, note that the MAPEs each represent one particular year at one particular location. We suspect that occasional deviant results like this are probably chance occurrences, rather than an indication of systematic failure of the MAV to do better than a naïve forecast at this location. More data would be required to confirm this hypothesis.

temperature forecasts are not as important for electricity producers in the West as they are elsewhere. Also, the consistency across nonwestern sites for the MAPE reduction is an encouraging suggestion that MAPE reductions in these parts of the country may not be highly sensitive to site choice.

The second thing to note in Table 1 is that, for the NWS forecast case, the load forecast errors for Florida (3.48) and Texas (3.98) are quite similar to those for the Hobbs et al. (1999) southern site (3.9); but, the load forecast errors for Vermont (2.01) and Ohio (3.36) are much lower than for the Hobbs et al.'s northeastern site (5.4). This is of interest because the Hobbs et al. study indicates that the benefits of MAPE reductions are smaller when the initial MAPE is smaller, and when the site is a northern site. Low MAPEs in this study's northern sites will tend to reduce the estimated benefits of NWS forecasts for these sites. However, because the benefits of MAPE reduction tend to be smaller anyway for northern sites, any downward bias from lower initial MAPEs for the northern sites will matter less because these are northern sites.

THE VALUE OF TEMPERATURE FORECASTS IN UNITED STATES ELECTRICAL GENERATION.

Having now established the connection between alternative temperature forecasts and electricity load forecast errors, the work of Hobbs et al. (1999) can be used to calculate cost savings from these alternative temperature forecasts. From this, estimates can be made of the overall national benefit of NWS forecasts, as well as the additional benefit that might be realized from further improvements relative to the NWS forecasts.

Hobbs et al. (1999) estimated four relationships (see Fig. 1) between the load forecast error and the economic cost of that error—one relationship for each of two regions (North and South) and one for each of two types of generating systems (Bard and Shaw). Hobbs et al.'s two relationships (Bard and Shaw) for their North region are used to estimate cost savings from alternative weather forecasts at our Vermont, Ohio, and Washington sites, and the results are then averaged over the Bard and Shaw systems at each site. Similarly, Hobbs et al.'s relationships for their South region are used at our Florida, Texas, and southern California sites, and the results are averaged over the Bard and Shaw systems. This procedure produces percentage cost increases (relative to MAPE = 0) for each of our six sites and each forecast. Last, cost increases for alternative forecasts are subtracted at each site to obtain site-specific cost reductions from better forecasts. These cost reductions are reported in Table 2.

Three steps are now required to estimate cost reductions for the United States as a whole from the data in Table 2. First, the site-specific results in Table 2 are averaged to get region-specific-estimated cost reductions for three U.S. regions identified as North, South, and West. Next, electricity generation data are obtained for these three regions, and an average production cost is applied to these generation amounts to estimate the generation costs by region. Finally, these regional generation costs are multiplied by the region-specific cost reductions from better forecasts, and these results are aggregated over regions to arrive at estimates of the U.S. national cost savings from better temperature forecasts. The details of these steps are explained below.

The site-specific cost reductions in Table 2 are averaged for the Vermont and Ohio sites to estimate cost reductions for our North region. Similarly, cost reductions in Table 2 are averaged for the Florida and Texas sites, and for the Washington and southern California sites, to estimate cost reductions for our South and West regions, respectively.

United States electricity generation data for our three regions are developed from Energy Information

Table 2. Percentage cost reductions for improved weather forecasts.			
	NWS vs persistence forecast	NWS vs MAV forecast	Perfect vs NWS forecast
Vermont	0.086	0.003	0.022
Ohio	0.092	0.003	0.061
Florida	0.488	0.037	0.180
Texas	0.592	0.055	0.277
Southern California	0.013	0.051	0.188
Washington	0.030	0.009	0.033

Administration data on electric utility net generation by census division and state for the year 2002. (see Energy Information Administration 2003).[7] For our north region, the following census regions are aggregated: New England, mid-Atlantic, East–North-central, and West–North-central. Similarly, for our South region, the South Atlantic, East–South-central, West–South-central, census regions are aggregated. The Mountain census region is split, with Arizona and New Mexico included in our South region, while the remaining states are included in our North region. The Pacific census region is also split with Alaska, included in the North, while the rest of the region (California, Oregon, Washington, and Hawaii) is aggregated into our West region. After these aggregations and divisions, about 1,367,000 million KWh (54%) of electricity generation is located in our South, 974,000 million KWh (38%) in our North, and 206,000 (8%) in our West regions.

For the average production cost of electricity, the number cited by Hobbs et al. (1999) as a typical cost—$20 per million watt hours—is used. This is equivalent to $20,000 per million killowatt hours (the units used for the electricity generation numbers above). Thus, the regional electricity generation amounts above are multiplied by $20,000 to produce the following estimated operating costs by region: $27.34 billion (South), $19.48 billion (North), and $4.13 billion (West).

Finally, these generation costs are multiplied by the estimated cost reductions through the use of better weather forecasts. This last set of calculations is summarized in Table 3.[8]

It is also possible to estimate the incremental benefits that are obtainable from improvements in forecasts starting from the current NWS forecast accuracy. For relatively small improvements in current forecast accuracy, this incremental benefit is estimated using the difference in benefits between the NWS forecast and the MAV forecast. This benefit increase is then divided by the percentage improvement in forecast accuracy between the NWS forecast

Table 3. Benefits of weather forecast improvements.					
		NWS vs persistence forecast		**Perfect vs NWS forecast**	
	Operating cost (million dollars)	Cost reduction (percent)	Benefits (million dollars per year)	Cost reduction (percent)	Benefits (million dollars per year)
North	19,478	0.0892	17.37	0.0419	8.17
South	27,341	0.5400	147.66	0.2286	62.50
West	4,127	0.0224	0.92	0.1105	4.56
Total			**165.95**		**75.23**

[7] Utility net generation is about two-thirds of total U.S. net generation. Nonutility net generation is excluded, because it is primarily for industrial use, and demand for it is likely to be less weather-sensitive. There is also likely to be less flexibility in the nonutility sector to commit units differently, depending on the weather forecast.

[8] The results in Table 3 are obviously point estimates, and it would be nice to know how sensitive they might be to changes in the data that are used to get them. To determine this fully would require a lot of data that are not available. However, it is possible to get some feel for the possible range around the point estimates in Table 3 by recalculating results alternately using the higher (lower) of Hobbs et al.'s (1999) Bard and Shaw system curves in Fig. 1, and the higher (lower) of our two site-specific cost savings estimates for each region. When these experiments are carried out, our results for the NWS versus persistence forecast swing up (down) by about $50 million per year (or 30%), while those for the perfect versus NWS forecast swing up (down) by about $30 million per year (or 40%).

and the MAV forecast to get the incremental benefit per percentage point improvement in accuracy. These incremental benefits are shown below in the last column of Table 4.

To value larger improvements in forecast accuracy, it is necessary to account for the fact that the benefits of forecast accuracy are a nonlinear function of accuracy. To account for this, a second-degree polynomial is fit to the four data points (perfect, NWS, MAV, and persistence) estimated for each region. Figure 2 shows the data points and fitted polynomials for the three regions. The fitted polynomial functions can then be used to estimate the percentage cost reductions in each region if the temperature error is reduced, for example, by one full degree Celsius starting from the current NWS error levels (the second set of points to the right of the origin on the horizontal axis). These cost reductions are then applied to the estimated costs to get the benefit of a 1°C improvement. The calculations, shown in Table 5, indicate that a 1°C improvement in forecast accuracy is worth about $59 million per year.[9]

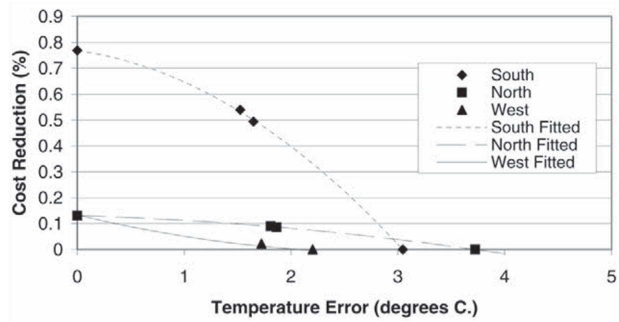

Fig. 2. Percentage cost reduction depends on error in temperature forecast.

DISCUSSION OF RESULTS. The NWS forecast produces a total U.S. benefit of $166 million per year, relative to the persistence forecast. Of this total benefit, $148 million is estimated to come from the South region. There are several reasons why most of the benefit comes from southern systems. First, a little more than half of the country's electricity is generated in the South region. Second, Hobbs et al.'s (1999) results indicate that the cost of a given load estimation error is roughly twice as high for their representative southern system as for their northeastern system. Third, the amount of load uncertainty (measured by MAPE with the NWS forecasts) is 39% higher for the South region than for the North, and 34% higher for the South than the West. Most importantly, however, the estimated reduction in MAPE from better temperature forecasts is much greater for the South region. On average, the MAPE reduction in going from a persistence forecast to the NWS forecast is almost twice as great in the South region as in the North (2.05 versus 1.08 percentage points), and almost 10 times as great in the South region as in the West (2.05 versus .21 percentage points). Presumably, this is due to the fact that air conditioning is a more important element of demand in the South, and temperature forecasts are primarily helpful in predicting electricity demanded for air conditioning.

The benefit from having a perfect forecast to replace the NWS forecast is estimated to be $75 million per year. This implies that about 70% of the total potential benefits of a perfect forecast versus a persistence forecast have already been realized by using the current NWS forecast. This is presumably due, in part, to the fact that the NWS forecasts are pretty good, and partly due to the fact that the incremental benefits of forecast improvements are highest for the initial improvements, and then decline as the forecast approaches perfection.

The incremental benefit of an improvement in forecast accuracy is estimated to be about $1.4 million per percentage point of improvement per year. For a larger 1°C improvement in accuracy, the benefit is about $59 million per year.

Table 4. Incremental benefits of one percentage point weather forecast improvement.

		NWS vs MAV forecast			
	Operating cost (million dollars)	Cost reduction (percent)	Benefits (million dollars per year)	Percentage forecast improvement	Benefit per percent (million dollars)
North	19,478	0.0030	.58	7	.082
South	27,341	0.0459	12.55	10.5	1.195
West	4,127	0.0299	1.24	11.4	.108
Total			**14.36**		**1.386**

[9] If this calculation is performed for a 1°F forecast accuracy improvement, the benefit is $37 million per year.

Table 5. Incremental benefits of 1°C weather forecast improvement.			
	Operating cost (million dollars)	Cost reduction (percent)	Benefits (million dollars per year)
North	19,478	0.0284	5.54
South	27,341	0.1883	51.48
West	4,127	0.0582	2.40
Total			**59.43**

Incremental benefits are relevant in assessing the merits of investments that will improve forecast accuracy. Sometimes investments that improve accuracy may be essentially one-time investments. In such cases, it would be appropriate to compare the cost of the investment to the present value of all future benefits from improved accuracy. These present value benefits would be $28 million for a 1% improvement in accuracy, and $1.2 billion for a 1°C improvement in accuracy, using a 5% discount rate over an infinite time horizon.[10]

In industries that are especially weather-sensitive, private meteorologists (or private meteorological services) are sometimes employed to produce forecasts tailored to the special requirements of the users in those industries. To decide whether paying for such privately provided forecasts makes sense, it would be appropriate to compare the cost of these forecasts to the economic benefit expected from the improvement in forecast accuracy. The results in Table 4 provide a rough estimate of the possible benefit per percentage point improvement in accuracy for electricity generators in the U.S. North, South, or West. To get a more accurate estimate, it would be necessary for an electricity generator to replicate the analysis presented here for the specific area that it serves and for the specific generating alternatives available to it.

Finally, it should be noted that the foregoing estimates only reflect the benefits of relatively short-term temperature forecasts used in scheduling units for next-day generation. There are additional electric utility benefits of weather forecasts that are not included in the above numbers, notably the value of multiday forecasts in scheduling maintenance. Thus, the estimated benefits from the temperature forecasts presented here certainly understate the total benefits realized by the electricity-generating industry from weather forecasts.

[10]The infinite horizon assumption makes sense if the improvement essentially lasts forever and any future maintenance or operating costs are also discounted and added to the initial investment outlay that produces the forecast improvement.

ACKNOWLEDGMENTS. We are grateful to Paul Dallavalle and Ben Hobbs for their assistance with this research. We also thank two anonymous referees and editorial board member Tom Fahey for their many helpful suggestions.

REFERENCES

Bard, J. F., 1988: Short-term scheduling of thermal-electric generation using Lagrangian relaxation. *Oper. Res.*, **36,** 756–766.

Dallavalle, J. P., M. C. Erickson, and J. C. Maloney III, 2004: Model output statistics (MOS) guidance for short-range projections. Preprints, *20th Conf. on Weather Analysis and Forecasting,* Seattle, WA, Amer. Meteor. Soc., CD-ROM, 6.1.

Energy Information Administration, 2003: Table 7: Electric utility net generation by census division and state. Electric Power Monthly, March, DOI/EIA-0226 (2003/03), 220 pp. [Available online at http://tonto.eia.doe.gov/FTPROOT/electricity/epm/02260303.pdf.]

Gunasekera, D., 2004: Economic issues relating to meteorological services provision. Bureau of Meteorology Research Centre, Research Rep. 102. [Available online at www.bom.gov/au/bmrc/pubs/researchreports/RR102.pdf.]

Hobbs, B. F., S. Jitprapaikulsarn, S. Konda, V. Chankong, K. A. Loparo, and D. J. Maratukulam, 1999: Analysis of the value for unit commitment decisions of improved load forecasts. *IEEE Trans. Power Syst.,* **14,** 1342–1348.

Khotanzad, A., R. C. Hwang, A. Abaye, and D. J. Maratukulam, 1995: An adaptive modular artificial neural network hourly load forecaster and its implementation in electric utilities. *IEEE Trans. Power Syst.,* **10,** 1716–1722.

——, R. Afkhami-Rohani, and D. Maratukulam, 1998: ANNSTLF—Artificial neural network short-term load forecaster-generation three. *IEEE. Trans. Power Syst.,* **13,** 1413–1422.

Shaw, J. J., 1995: A direct method for security-constrainted unit commitment. *IEEE Trans. Power Syst.,* **10,** 1329–1342.